

AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application.

1. (Currently Amended) A method for allocating resources to a plurality of applications, comprising:
determining, by a computer, available resources of a networked computing system, wherein the available resources comprise processing resources, networking resources, and storage resources;
determining, by the computer, for each application, required resources of the application, wherein the required resources of each application specify a type of servers on which the application can be hosted, the type of servers being based on attributes of the servers including attributes relating to processor architecture and processing power, wherein the servers are part of the processing resources;
determining, by the computer, an assigned subset of the available resources for each application as a function of the required resources of the application and the available resources, wherein determining the assigned subset of available resources for each application is based on a linearized objective function that reduces communication delays between resources of the subset of the available resources in conformance with bandwidth capacity requirements of the application and in conformance with network bandwidth limitations, wherein the linearized objective function includes a linear combination of variables; and
associating the applications with the assigned subsets of resources.

2.-3. (Cancelled)

4. (Previously Presented) The method of claim 1, wherein the networking resources comprise network switches.

5. (Cancelled)

6. (Previously Presented) The method of claim 1, wherein the storage resources comprise a storage area network, wherein the storage area network includes at least one pair of redundant core switches coupled to storage devices, the core switches coupled to the processing resources via a plurality of edge switches.

7. (Currently Amended) The method of claim [[6]]1, wherein ~~assignment of the determining~~ the assigned subset of available resources for each application is performed by solving a mixed-integer programming problem.

8. (Previously Presented) The method of claim 7, wherein the available resources include network switches, and the mixed-integer programming problem reduces communication delays between resources of each subset of the available resources by reducing data traffic on network links that interconnect the network switches.

9. (Currently Amended) A system comprising:
a computer having a processor and configured to:
~~means for determining~~ determine available resources of a networked computing system, wherein the available resources comprise servers, networking resources, and storage resources;
~~means for determining~~ determine required resources for each application of a plurality of applications, wherein the required resources for each application specify a type of the servers on which the application can be hosted, the type of servers being based on attributes relating to processing architecture and processing power;
~~means for determining~~ determine an assigned subset of the available resources for each application as a function of the required resources of the application and the available resources, wherein determining the assigned subset of available resources for each application is based on a linearized objective function that reduces communication delays between resources of the subset of the available resources in conformance with bandwidth capacity requirements of the application and in conformance with network bandwidth limitations, wherein the linearized objective function includes a linear combination of variables; and
~~means for associating~~ associate the applications with the assigned subsets of resources.

1 10.-11. (Cancelled)

1 12. (Currently Amended) The computer-readable medium of claim [[10]]26, wherein the
2 processing resources comprise servers each having at least one processor.

1 13. (Currently Amended) The computer-readable medium of claim [[10]]26, wherein the
2 networking resources comprise network switches.

1 14. (Cancelled).

1 15. (Currently Amended) The computer-readable medium of claim [[10]]26, wherein the
2 storage resources comprise a storage area network, wherein the storage area network includes at
3 least one pair of redundant core switches coupled to storage devices, the core switches coupled to
4 the processing resources via a plurality of edge switches.

1 16. (Previously Presented) The computer-readable medium of claim 15, wherein reducing the
2 communication delays between resources comprises solving a mixed-integer programming
3 problem.

1 17. (Previously Presented) The computer-readable medium of claim 16, wherein the
2 available resources include network switches coupled with the processing resources, and the
3 mixed-integer programming problem reduces communication delays between resources of each
4 subset of the available resources by reducing data traffic on network links that interconnect the
5 network switches.

1 18. (Currently Amended) A system, comprising:
2 a plurality of network-coupled processing resources;
3 a plurality of storage resources network-coupled to the processing resources, wherein the
4 processing and storage resources are allocated to a plurality of applications;
5 a computing arrangement configured to:
6 determine, for each application of the plurality of applications, required resources
7 of the application, wherein the required resources of each application specify a type of servers on
8 which the application can be hosted, the type of servers being based on attributes of the servers
9 including attributes relating to processor architecture and processing power, wherein the servers
10 are part of the processing resources;
11 determining-determine an assigned subset of the processing and storage resources
12 for each application as a function of the required resources of the application and the processing
13 and storage resources, wherein determining the assigned subset of processing and storage
14 resources for each application is based on a linearized objective function that reduces
15 communication delays between resources of the subset of the processing and storage resources in
16 conformance with bandwidth capacity requirements of the application and in conformance with
17 network bandwidth limitations, wherein the linearized objective function includes a linear
18 combination of variables; and
19 associate the applications with the assigned subsets of processing and storage
20 resources.

1 19.-20. (Cancelled)

1 21. (Original) The system of claim 18, wherein the storage resources comprise a storage area
2 network, wherein the storage area network includes at least one pair of redundant core switches
3 coupled to storage devices, the core switches coupled to the network via a plurality of edge
4 switches.

1 22. (Previously Presented) The system of claim 21, wherein the computing arrangement is
2 configured to determine each assigned subset by solving a mixed-integer programming problem.

23. (Currently Amended) The system of claim 22, wherein the processing resources are coupled by network switches, and the mixed-integer programming problem reduces communication delays between resources by reducing data traffic on network links that interconnect the network switches.

24. (Currently Amended) The method of claim 1, wherein the ~~required resources of each application is specified in resource requirements that include attributes of the processing resources, wherein the attributes relating to processor architecture and processing power~~ specify processor type and processor speed.

25. (Currently Amended) The method of claim 24, wherein the resource requirements further specify storage patterns of files for each application, wherein determining the assigned subset is based on the resource requirements that further specify storage patterns of files.

26. (Currently Amended) ~~The computer-readable medium of claim 10~~ A computer-readable medium configured with instructions for causing a processor of a computer to allocate resources to a plurality of applications, comprising:

determining available resources of a networked computing system, wherein the available resources comprise processing resources, networking resources, and storage resources;

determining, for each application, required resources of the application;

determining an assigned subset of the available resources for each application as a function of the required resources of the application and the available resources, wherein determining the assigned subset of available resources for each application is based on a

linearized objective function that reduces communication delays between resources of the subset of the available resources in conformance with bandwidth capacity requirements of the application and in conformance with network bandwidth limitations, wherein the linearized objective function includes a linear combination of variables; and

associating the applications with the assigned subsets of resources,

wherein the required resources of each application is specified in resource requirements that include attributes of the processing resources, wherein the attributes specify processor type and processor speed.

27. (Currently Amended) The computer-readable medium of claim 26, wherein the resource requirements further specify storage patterns of files for each application, wherein determining the assigned subset is based on the resource requirements that further specify storage patterns of files.

28. (Previously Presented) The method of claim 1, wherein the network bandwidth limitations are expressed as linear constraints, and wherein determining the assigned subset of the available resources for each application is a linear optimization problem.

29. (Previously Presented) The method of claim 28, wherein determining the assigned subset of the available resources for each application is a mixed integer programming problem.

30. (Previously Presented) The method of claim 1, wherein determining the assigned subset of the available resources for each application is based on the linearized objective function to reduce a number of hops between processing resources in the assigned subset.

31. (Previously Presented) The method of claim 1, wherein the linearized objective function is provided by substituting products of binary variables in a non-linear objective function with replacement binary variables in the linearized objective function.

32. (Currently Amended) The computer-readable media of claim [[10]]26, wherein the network bandwidth limitations are expressed as linear constraints, and wherein determining the assigned subset of the available resources for each application is a linear optimization problem.

33. (Previously Presented) The computer-readable media of claim 32, wherein determining the assigned subset of the available resources for each application is a mixed integer programming problem.

1 34. (Currently Amended) The computer-readable media of claim [[10]]26, wherein
2 determining the assigned subset of the available resources for each application is based on the
3 linearized objective function to reduce a number of hops between processing resources in the
4 assigned subset.

1 35. (Currently Amended) The computer-readable media of claim [[10]]26, wherein the
2 linearized objective function is provided by substituting products of binary variables in a non-
3 linear objective function with replacement binary variables in the linearized objective function.